

Crosstalk, Part 2

Simulating Crosstalk Effects

Douglas Brooks, President
UltraCAD Design, Inc.
April 2004

ABSTRACT

It is known that forward crosstalk increases (for all practical purposes) with increasing coupled length, but has a pulse width that is constant. Backwards crosstalk, on the other hand, rises quickly (within the critical region) to a constant maximum, but has a pulse width that increases with increasing coupled length. Simulations using HyperLynx® LineSim® show this very effectively and clearly. The tool can illustrate how impedance loading of the victim trace can impact the magnitude, and even the polarity, of the backward crosstalk pulse. HyperLynx also can be used to clearly illustrate how the backward crosstalk pulse is twice the propagation time through the coupled region plus one rise time, how the crosstalk signal is impacted by the relationship between the traces and their reference planes and also with each other, and how an aggressor AC signal's period can interact with the length of the coupled region to create some surprising crosstalk effects.

INTRODUCTION

In a previous article¹ I talked about the nature of crosstalk, especially forward vs backward crosstalk and how forward crosstalk is virtually zero in a stripline environment. This article will introduce several HyperLynx simulations illustrating various characteristics of crosstalk. The reader will gain an understanding of (a) how to use HyperLynx LineSim to simulate crosstalk, and (b) how to recognize various crosstalk effects.

We will illustrate a variety of simulations:

- Basic crosstalk concepts
- Difference between microstrip and stripline crosstalk effects
- The effect of trace terminations on crosstalk
- Crosstalk sensitivity to trace height and separation
- What happens when there is coupling over only part of the trace length

Another topic, crosstalk coupling when we tune traces with closely coupled "serpentine" traces, has been covered in a separate article.²

BASIC MODEL

To simulate basic crosstalk concepts without any spurious effects getting in the way, we will structure a fairly specific, simple model. This model will be designed with the intent of excluding all spurious effects except those caused by crosstalk between the traces themselves. The stripline portion of the model is shown in Figure 1.

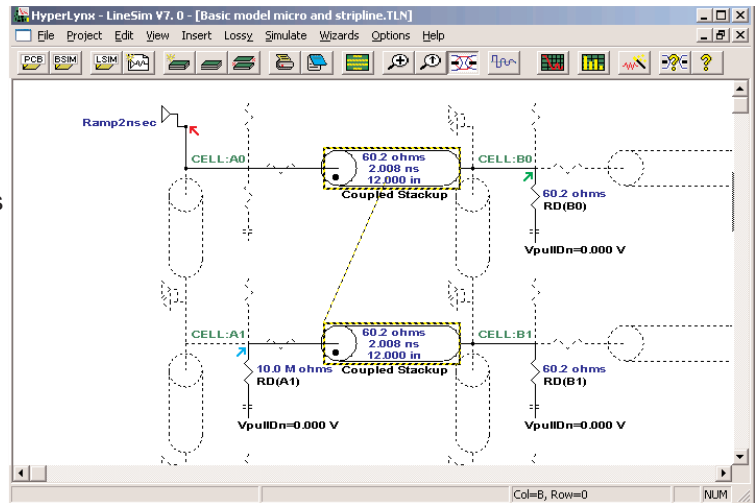


Figure 1 - Basic HyperLynx LinSim model for simulation.

The stackup for this model is shown in Figure 2. In HyperLynx, the stackup is created or edited by the menu selection Edit/Stackup. In this paper we are going to be using the layers labeled "Top" (microstrip) and "Stripline" (Signal layer 4). The stripline stackup values have been carefully created to achieve a propagation speed of almost exactly 6"/ns for ease in understanding the results.

Layer Name	Type	Usage	Thickness mils, oz	Er	Test Width mils	Z0 ohm
1 TOP	Metal	Signal	0.504	<Auto>	8	56
2	Dielectric	Substrate	5	3.9		
3 VCC	Metal	Solid Plane	1.009	<Auto>		
4	Dielectric	Substrate	10	3.9		
5 LayerNN001	Metal	Signal	0.504	<Auto>	6	68.6
6	Dielectric	Substrate	10	3.9		
7 Stripline	Metal	Signal	0.504	<Auto>	8	61.3
8	Dielectric	Substrate	10	3.9		
9 GND	Metal	Solid Plane	1.009	<Auto>		
10	Dielectric	Substrate	23	3.9		
11 BOTTOM	Metal	Signal	0.504	<Auto>	8	111

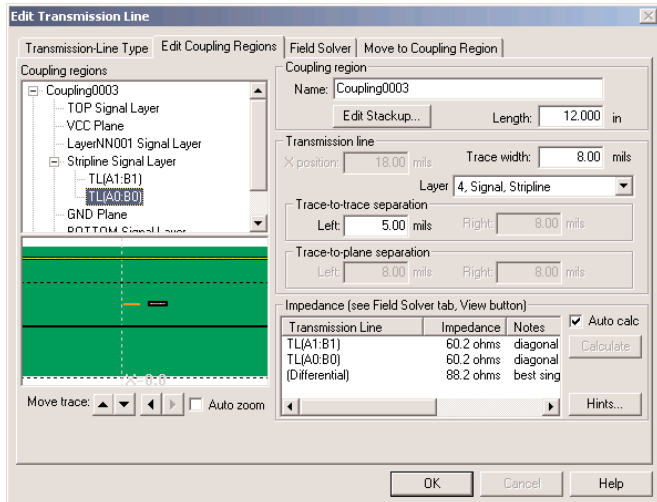
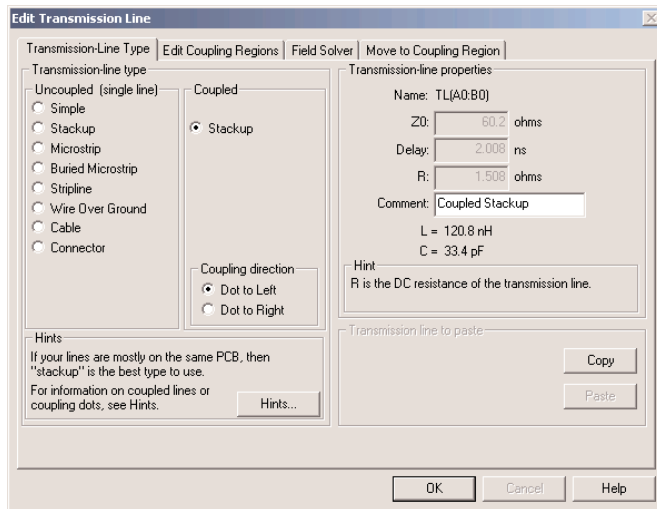
Measurement units: English
Metal thickness as: Weight

3D Stackup Diagram (mils, oz):
 TOP: 5 mils (0.504 oz)
 VCC: 1.009 oz (1.009 mils)
 LayerNN001: 10 mils (0.504 oz)
 Stripline: 10 mils (0.504 oz)
 GND: 1.009 oz (1.009 mils)
 BOTTOM: 23 mils (0.504 oz)

Total thickness: 63.4 mils
No errors in stackup.

Figure 2 - HyperLynx Stackup Editor.

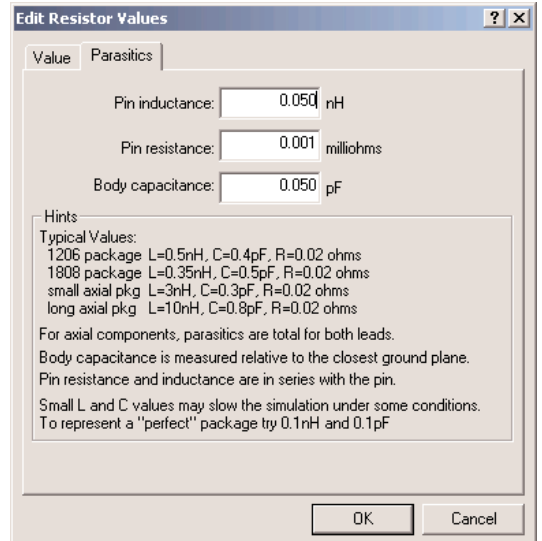
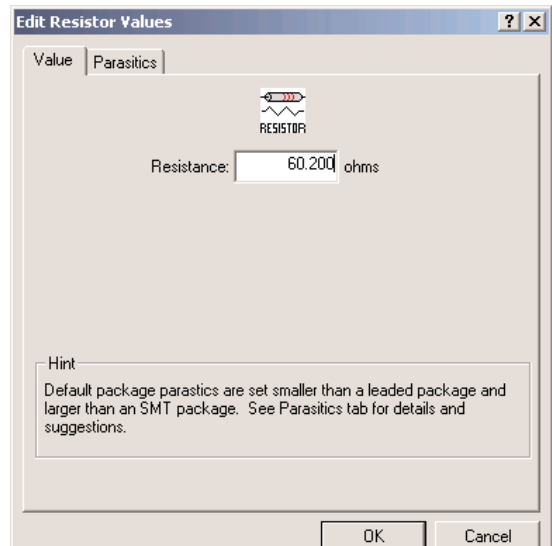
If you right-click one of the transmission line segments the "Edit Transmission Line" window displays (Figure 3a). The most significant entry here is that the "Coupled Stackup" radio button has been selected. The "Edit Coupling Regions" tab (Figure 3b) reveals the name and the layer of the coupling region. But more importantly, there are places to edit the length and width of the trace and the separation from the adjacent (coupled) trace. This illustration shows a length of 12 inches, a trace width of eight mils, and a separation of only five mils. We can expect the crosstalk to be fairly significant for two traces ten mils away from their reference plane (Figure 2) and only five mils apart (Figure 3b).



Figures 3a and 3b - Transmission Line Edit windows.

We need terminating resistors for the traces. If you select a resistor "holder" and then right-click it, the Edit Resistor Values window displays (Figures 4a and b). We have set the values of the two far end resistors to

60.2 ohms, to properly terminate the transmission lines. The near resistor on the victim line has been set to ten M ohms to simulate an open circuit or a high impedance pin on a device. Thus, the backward crosstalk signal on the victim line will reflect off this point (with a 100% reflection coefficient³) and travel back toward the far end of the victim line.



Figures 4a and 4b - Setting terminating resistor values (top) and parasitics (bottom).

HyperLynx allows you to set parasitic values for the passive devices (Figure 4b). This can be useful or desirable if you want to try to precisely simulate how a circuit will work with "real world" devices. For this theoretical basic model, I have set all the parasitic values to the minimum allowed by HyperLynx. This is so we can look as closely as possible to the crosstalk effects, isolated from any other effects that might distort the results.

Finally, we need a driver to drive the aggressor trace. That is identified in Figure 1 as "Ramp2nsec." Because there are some things I specifically want to identify from the basic results, I have created a driver with some very special characteristics. This driver rises almost perfectly linearly from zero to full value in two ns. That compares with a typical device that would have normal "roll-on" and "roll-off" characteristics. The purpose for this will become clear later. But the driving signal in this model is either constant or it is linear at 50%/ns. There is virtually no roll-on or roll-off.

To create a driver like this, start by Selecting Edit/Databook IC Models (.Mod). Figure 5a shows the Edit .MOD Model for the Device Model TDR in the GENERIC MOD Library. Figure 5b shows how I edited this model, and then saved it as Ramp2nsec in the USER.mod library for my use. The edited model is the driver we are going to use in this simulation.

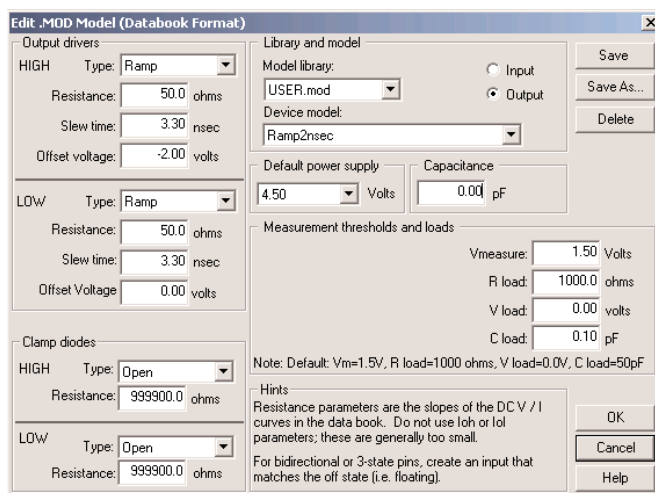
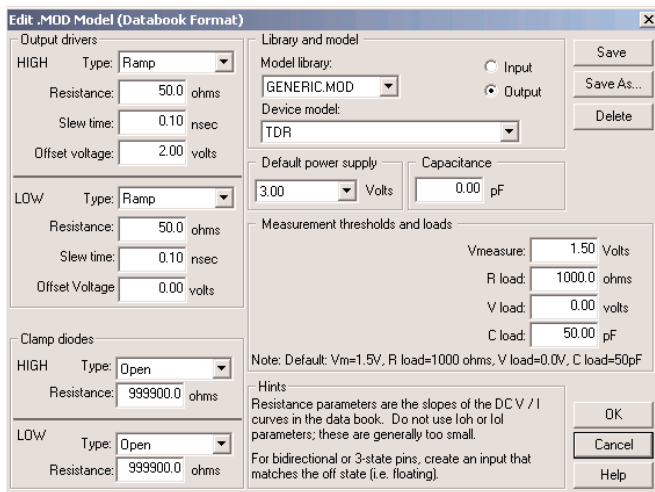


Figure 5a and 5b - Device model windows. Figure 5a (top) shows a standard GENERIC MOD model for a TDR. Figure 5b (bottom) shows how this model was edited to create Ramp2nsec device model used in the simulation.

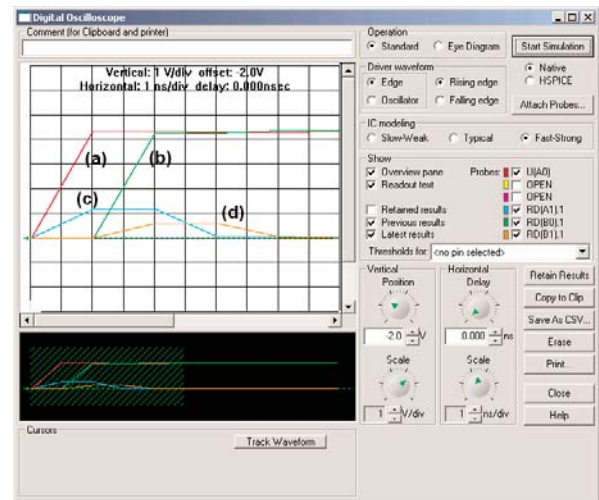


Figure 6 - Result from the basic simulation.

INITIAL RESULTS

Running the simulation (Simulate/Run Scope..) leads to the results shown in Figure 6. The red trace (a) in Figure 6 is the driver. Note how it rises from zero to full value (about 4.4 Volts) in almost exactly two ns. The green trace (b) is the driven signal showing up at the far end of the driven trace two ns later (the trace is 12 inches, or two ns long.) The blue trace (c) is the backward crosstalk signal arriving at, and reflecting from, the ten Mohm resistor at the front of the victim trace. In this example, the magnitude of the signal here is approximately 1.2 Volts.⁴ The orange trace (d) is this reflected signal showing up at the far end of the victim trace, two ns later. Its magnitude is about .6 V. The near end backward crosstalk pulse is actually .6 V, too, but it shows up in the model as twice that value because of the 100% positive reflection at the end of the (effectively) unterminated transmission line.

It is instructive to look at pulse widths and shapes here. And now it begins to come clear why we structured such a special driver. The crosstalk signal has the same rise time as the driver does (two ns). And since there is no perceptible roll-on or roll-off, it is very clear where these signals start and finish.

In the article referenced in Footnote 1, I pointed out that the backward pulse width is twice as wide as the propagation time through the coupled region. That is approximately true if you measure it from the midpoint of the rise and fall times. What is precisely true (and can be seen in this special case) is that the backward pulse width is twice the propagation time through the coupled region PLUS one rise time! Note that the rise time of the driving pulse is exactly two ns. The

coupled region is two ns long; twice that is four ns. The backward crosstalk pulse in Figure 6 is six ns from start to finish-twice the coupled region (4) plus one rise time (2).

The reason is this. As the very first part of the driven signal enters the coupled region, the backward crosstalk signal begins to form. As the very first part of the driven signal arrives at the far end of the coupled region, it generates a backward pulse component that is twice the coupled region behind the first. It is the length of the coupled region later, and it is the length of the coupled region further away from the beginning. These combine to double the coupled region width. But that only concerns the very first part of the driver signal (the beginning of the rise time). The last part of the driven signal (the top of the rise time) also contributes to the backward crosstalk pulse. It is one rise time later than the first part of the driven signal. Altogether, these factors combine to a backward pulse that is twice as wide as the coupled region plus one rise time.

COUPLED LENGTH

Figure 7 illustrates what happens as we change the coupled length. The driver is as before. The longest backward crosstalk pulse is the same as in Figure 6. Two other pulses are shown: one each for a six inch coupled region (one ns) and a three inch coupled region (.5 ns). Each of these backward pulses has a width of twice the coupled region plus one rise time. When the coupled region is the "critical length,"⁵ i.e. one-half the rise time, the signal just has time to rise to its maximum value and then decay again. So at the

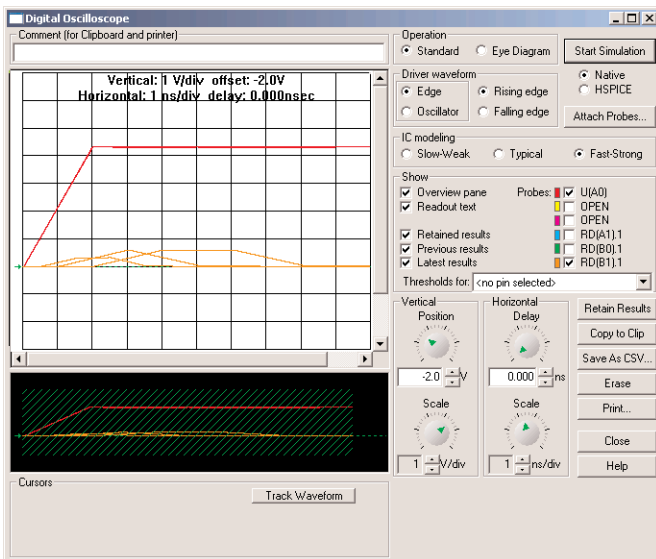
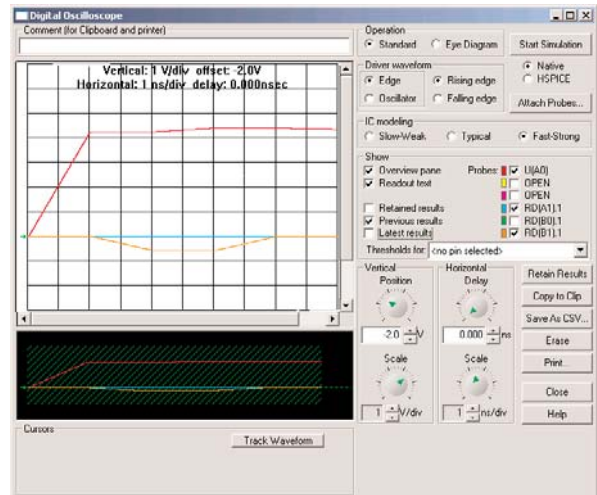
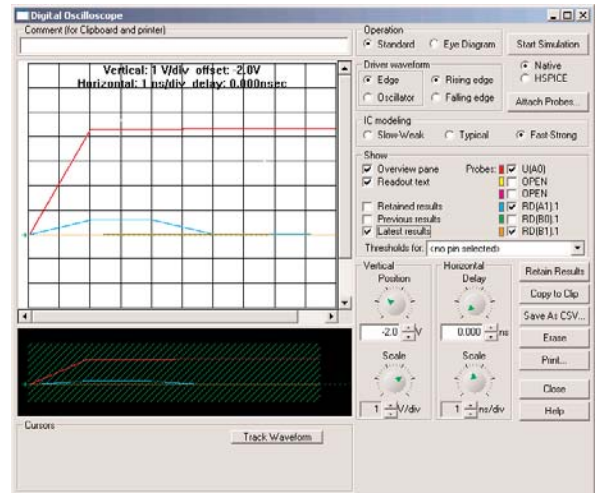


Figure 7 - Simulation results for three coupled lengths, 3" (a), 6" (b), and 12" (c).

critical length the backward pulse width is a triangle. If the coupled region is shorter than the critical length, the width follows the same rule as before, but the crosstalk signal does not have time to rise to its maximum value.

This, then, illustrates the fundamental rule: *The backward pulse increases in amplitude with the length of the coupled region until the critical length, then it rises no further. The pulse width is twice the coupled region plus one rise time.*

TERMINATIONS



Figures 8a (top) and 8b (bottom) - The impedance at the near-end of the victim trace can have a dramatic effect on the magnitude, and even the polarity, of the backward crosstalk pulse.

It is interesting to see what happens as we vary the terminating resistor at the front (near end) of the victim line. If the victim trace is terminated in its characteristic impedance at the near end, the backward crosstalk pulse travels back to it (Figure 8a) but then is completely dissipated at that point. There is no forward reflection. If, on the other hand, the trace is shorted at the near end, Figure 8b, then there is a

100% negative reflection of the backward pulse. Note that the pulse at the far end of the victim trace in Figure 8b is exactly the same shape and amplitude as the pulse in Figure 6, but opposite in sign. (There is no signal at the near end of the victim trace in Figure 8b, of course, because the trace is shorted at that point!)

Clearly, the magnitude (and even the polarity) of the crosstalk signal depends on the termination at the near end of the victim line. Any devices that exist on that portion of the line (including another device driver at the beginning of the victim line) contribute to that loading. And those devices may present different loading characteristics depending on their state (i.e. "on" or "off" or "high" or "low"). This makes estimating and troubleshooting crosstalk issues that much tougher.

SEPARATION AND HEIGHT ABOVE PLANE

The two primary ways to reduce crosstalk are well known:⁶ route traces close to their reference plane, and spread the traces further apart. Figure 9 simulates each of these effects. If the trace separation increases from five mils to ten mils, the magnitude of the backward crosstalk signal reduces as shown in Figure 9a. If the reference plane is brought closer to the traces (from ten mils to five mils) the magnitude also reduces, as shown in Figure 9b. It should be noted that when the trace is placed closer to the reference plane, the characteristic impedance drops. Therefore, the terminating resistors were adjusted accordingly.

MORE TYPICAL DRIVER

All the simulations so far have been with the special ramp driver described in Figure 5b. There are some slight changes when we do the simulation using a more typical high-speed device model. Figure 10 illustrates the model with a driver selected from the EASY.MOD library, the CMOS 3.3V ULTRAFast device. Figure 11 illustrates the results of this model for the same stripline trace configuration used in Figure 6. As is apparent, the results are very similar.

It gets more interesting when we do this simulation on the microstrip (top) layer (Figure 12). The stackup is as shown in Figure 2. The trace impedance has changed because the geometry is now

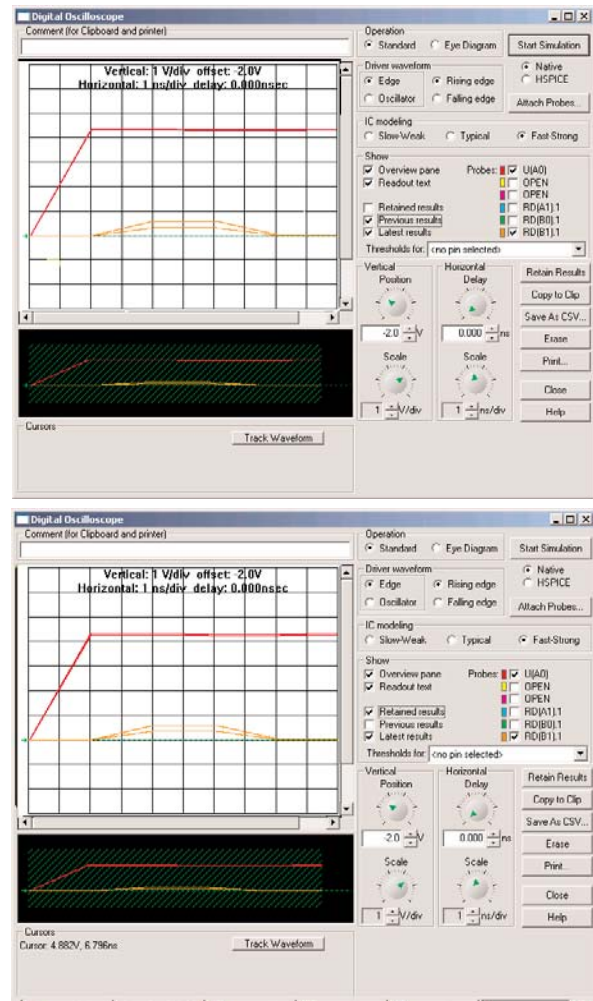


Figure 9a and 9b - The magnitude of the crosstalk pulse decreases with increasing separation, Figure 9a (top), and with decreasing height above the underlying plane, Figure 9b (bottom).

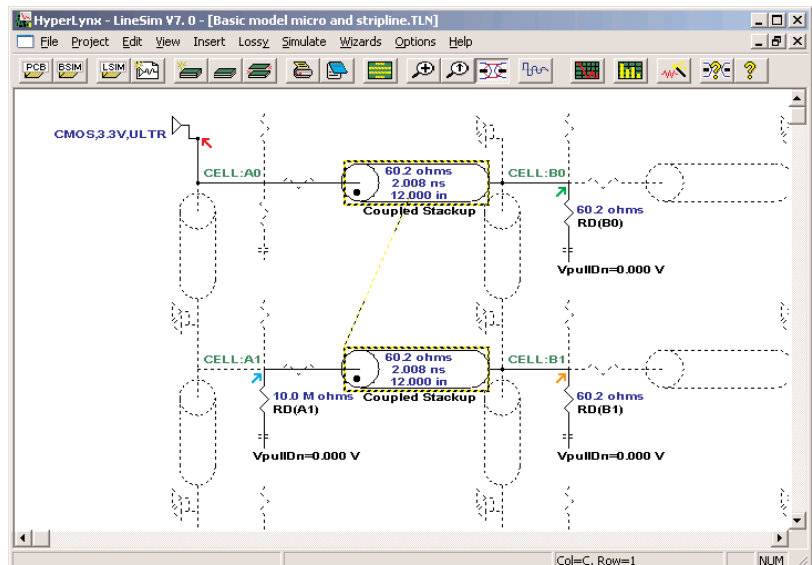


Figure 10 - Stripline simulation with a more conventional driver model.

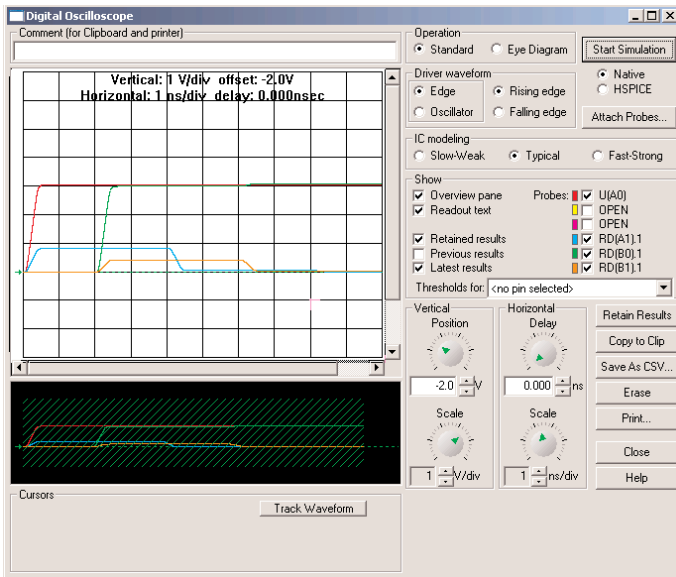


Figure 11 - Simulation result of the stripline model of Figure 10.

microstrip instead of stripline, and the terminating resistors have therefore been adjusted. Similarly, the propagation speed has changed.⁷ The simulation result is as shown in Figure 13.

The simulation is similar to that in Figure 11 except for the sharp negative-going pulse at the far end of the victim trace. This is the forward crosstalk pulse. Recall from the article referenced in Footnote 1, there are two coupling effects that take place in a crosstalk situation, capacitive and inductive. Capacitive coupling sends signal components in both directions in the victim trace, but inductive coupling only sends a signal component in the backward direction. In a stripline environment, the two components almost perfectly cancel in the forward direction, but reinforce each other in the

backward direction. Thus, there is a well-defined backward crosstalk signal but no forward crosstalk signal. But in a microstrip environment, the two forward components do not perfectly cancel. Thus there is a net forward crosstalk signal that can appear in microstrip environments that almost never occurs in stripline environments.

IMPACT OF COUPLED LENGTH

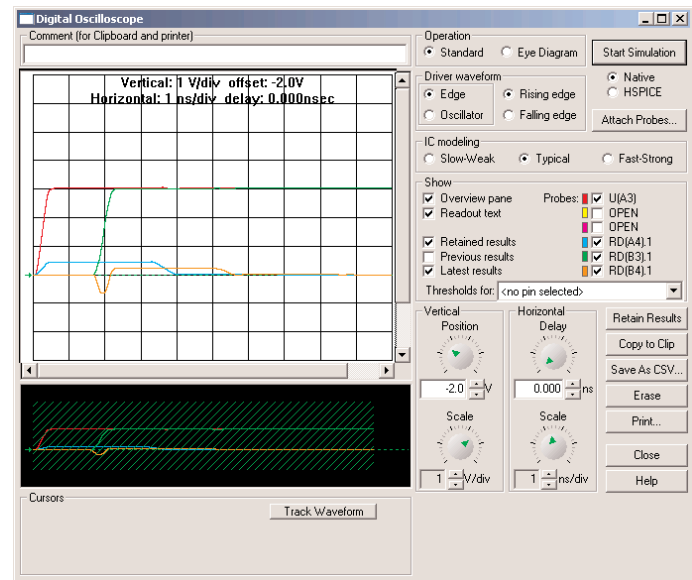


Figure 13 - Simulation result of the model in Figure 12.

The article referenced in Footnote 1 also points out that

the backward crosstalk signal is relatively constant in magnitude and has a pulse width twice as wide as the propagation time down the coupled region, but the forward crosstalk signal has an amplitude that continually increases as the coupled region increases and has a fixed pulse width equal to the rise time (or the fall time) of the aggressor signal.⁸

This result is illustrated in Figure 14, where three simulations are shown. The simulations are identical except for the length of the coupled region (six, 12 and 24 inches respectively.) Although forward crosstalk is rarely a problem in typical PCB designs (because the coupled length is

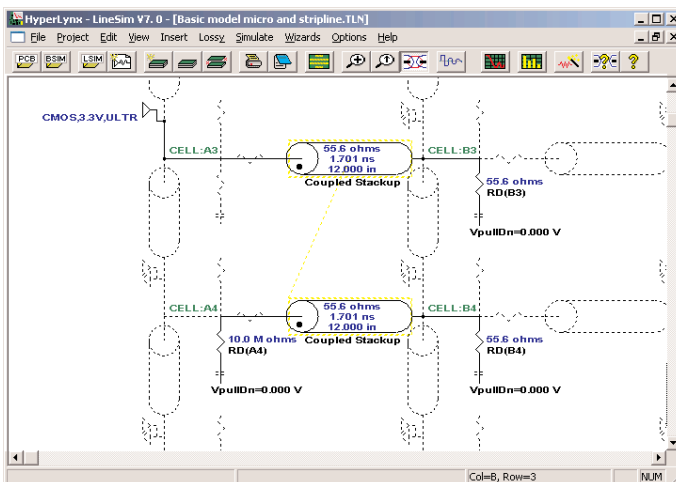


Figure 12 - Microstrip model with a more typical driver model.

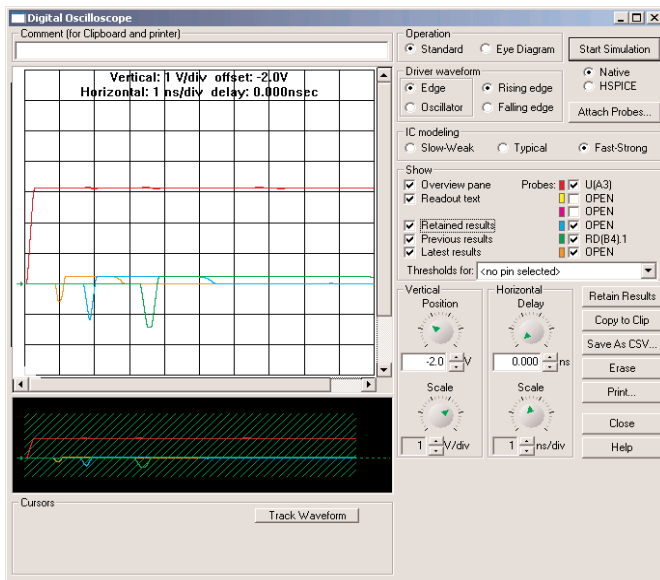
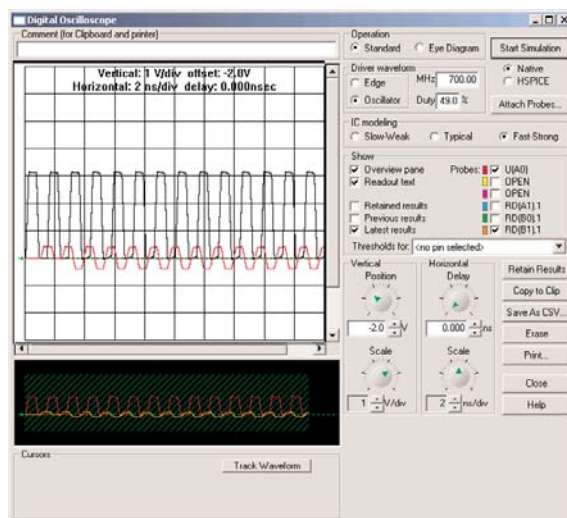
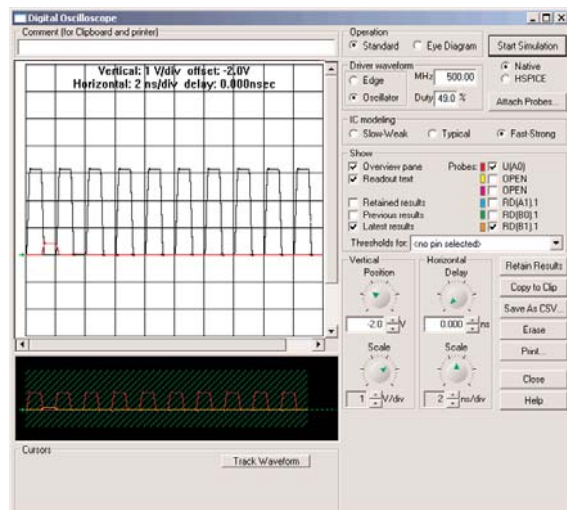
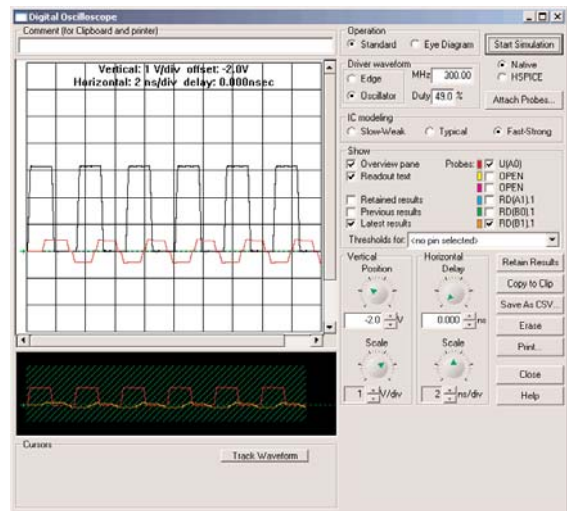


Figure 14 - Simulation results when the coupled length is varied from 6 to 12 to 24 inches.

rarely long enough to generate a significant signal), many designers follow the rule that if crosstalk is a performance issue or concern, all crosstalk sensitive traces will be routed in a stripline environment. This should be considered a conservative rule. Many designers have no particular crosstalk problems when routing in either environment.

OSCILLATING (AC) DRIVER

Up to now we have looked at simulations with a single pulse driven down the trace. Some surprising things can happen if we look at the more typical case of a clock signal driven down the trace. Consider the model shown in Figure 10 with the trace length adjusted to six inches (one ns). Let's drive this model with three different square wave frequencies, 300 MHz, 500 MHz, and 700 MHz, respectively. The results are shown in Figures 15a, 15b, and 15c, respectively. It is striking that there is no apparent crosstalk signal when the clock rate is 500 MHz. That does not mean that there is no crosstalk! The clock half-cycle is one ns. At time equal one ns there is a clock transition going in one direction. One ns later there is a clock transition going in the other direction. One ns later....., etc. The time of these transitions happens to coincide exactly with the trace length (one ns). Thus, the crosstalk components caused by each transition are canceling out. Imagine the difficulty of troubleshooting a crosstalk problem when varying frequencies are involved!



Figures 15a, 15b, and 15c - Crosstalk effects can be an unexpected function of frequency. Here are the results from the same model driven by an aggressor signal of 300 MHz (a, top), 500 MHz (b, middle), and 700 MHz (c, bottom). The larger waveform is the driver. The smaller one is the resulting crosstalk signal.

COUPLED OVER PARTIAL LENGTH

As a final simulation illustration, consider the case of two traces which are coupled over only part of their length. Figure 16 illustrates such a model. These are simulated as microstrip traces to incorporate the added complexity of a forward crosstalk pulse. The entire propagation delay is nine inches (1,281.4 ns) along the aggressor trace and 16 inches (2,287.3 ns) along the victim trace. The result of this simulation is shown in Figure 17. Note in particular that the forward and backward crosstalk signals arrive at the far end of the victim trace at significantly different times.

If you were an engineer and saw these traces on a scope for a signal trace and for a (supposedly) quiet adjacent trace, it would not be readily intuitive that these signals were related! That is one of the difficulties in troubleshooting crosstalk problems if you

are not familiar with what they look like or if you are not sensitive to the possibility that they could even exist.

SUMMARY

HyperLynx LineSim is a very effective tool for simulating crosstalk. It is equally effective for simulating actual circuits and for simulating hypothetical situations to see what the effects might look like. Crosstalk effects are quite similar in microstrip and in stripline environments, except forward crosstalk can become an issue in the former. Simulations can illustrate quite clearly the fundamental rule that forward crosstalk grows continuously (for all practical purposes) with increasing coupled length, but its pulse width remains constant. Backward crosstalk, on the other hand, rises quickly (within the critical region) to a constant maximum amplitude, but has a pulse width that increases with increasing coupled

length. The maximum amplitude, and even its polarity, can be a function of impedance loading at the front (near end) of the victim trace. Using a controlled driver waveform, it is easy to show that the backward crosstalk pulse width is twice the propagation time through the coupled region plus one rise time. Finally, for AC waveforms, there can be an interaction between the period of the waveform and the propagation time through the coupled region that can have a very significant effect on the resulting crosstalk signal.

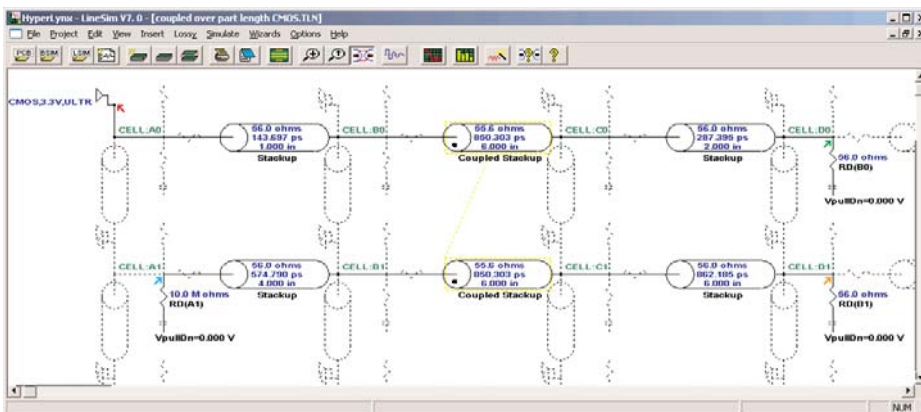


Figure 16 - Simulation of a condition where the traces are coupled over only a portion of their length.

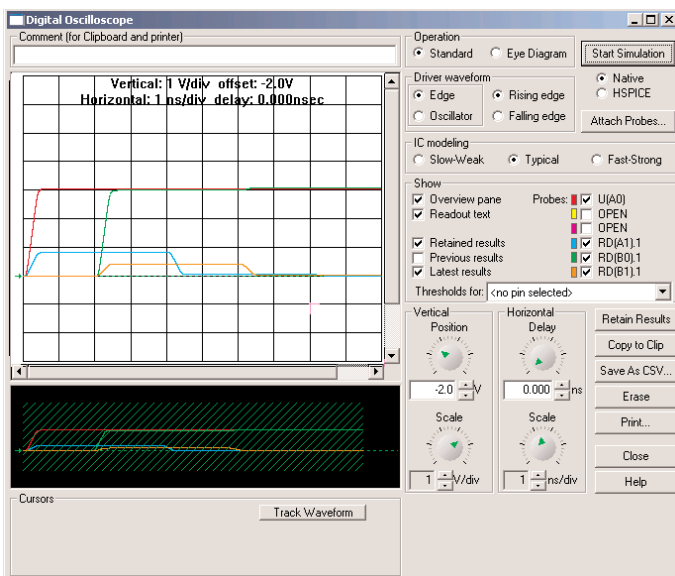


Figure 17 - Simulation result of the model in Figure 16.

FOOTNOTES

1. "Crosstalk, Part 1; Understanding Forward vs Backward," available at <http://www.mentor.com/pcb/techpapers/>.
2. "Adjusting Signal Timing, Part 2; Crosstalk Effects in Serpentine Traces," available at <http://www.mentor.com/pcb/techpapers/>.
3. Information on reflection coefficients can be found in a great many sources. See for example, "Transmission Line Terminations; It's the End that Counts," available at <http://www.mentor.com/pcb/techpapers/>.

4. This crosstalk signal is quite high because we have structured an example with very strong coupling (traces relatively far from their reference plane and close to each other), and because there is a 100% reflection here at the effectively unterminated end of the transmission line.

5. This is the same critical length we talk about with transmission lines. There are many, many articles on this topic. See, for example, "Propagation Times and Critical Length; How They Interrelate," available at <http://www.mentor.com/pcb/techpapers/>.

6. Brooks, Douglas, Signal Integrity Issues and Printed Circuit Board Design, Prentice Hall, 2003, Chapter 12.

7. "Microstrip Propagation Times; Slower than We Think," available at <http://www.mentor.com/pcb/techpapers/>.

8. Same as Footnote 1.

to those without a technical background. You can visit his web page at <http://www.ultracad.com> and e-mail him at doug@ultracad.net.

ABOUT THE AUTHOR

Douglas Brooks has a B.S and an M.S in Electrical Engineering from Stanford University and a PhD from the University of Washington. During his career has held positions in engineering, marketing, and general management with such companies as Hughes Aircraft, Texas Instruments and ELDEC.

Brooks has owned his own manufacturing company and he formed UltraCAD Design Inc. in 1992. UltraCAD is a service bureau in Bellevue, WA, that specializes in large, complex, high density, high speed designs, primarily in the video and data processing industries. Brooks has written numerous articles through the years, including articles and a column for Printed Circuit Design magazine, and has been a frequent seminar leader at PCB Design Conferences. His primary objective in his speaking and writing has been to make complex issues easily understandable

For more information, call us or visit: www.mentor.com/pcb

Copyright © 2004 Mentor Graphics Corporation. This document contains information that is proprietary to Mentor Graphics Corporation and may be duplicated in whole or in part by the original recipient for internal business purposes only, provided that this entire notice appears in all copies. In accepting this document, the recipient agrees to make every reasonable effort to prevent the unauthorized use of this information. Mentor Graphics is a registered trademark of Mentor Graphics Corporation. All other trademarks are the property of their respective owners.

Corporate Headquarters
Mentor Graphics Corporation
8005 S.W. Boeckman Road
Wilsonville, Oregon 97070 USA
Phone: 503-685-7000

Silicon Valley Headquarters
Mentor Graphics Corporation
1001 Ridder Park Drive
San Jose, California 95131 USA
Phone: 408-436-1500
Fax: 408-436-1501

Europe Headquarters
Mentor Graphics Corporation
Deutschland GmbH
Arnulfstrasse 201
80634 Munich
Germany
Phone: +49.89.57096.0
Fax: +49.89.57096.40

Pacific Rim Headquarters
Mentor Graphics (Taiwan)
Room 1603, 16F,
International Trade Building
No. 333, Section 1, Keelung Road
Taipei, Taiwan, ROC
Phone: 886-2-27576020
Fax: 886-2-27576027

Japan Headquarters
Mentor Graphics Japan Co., Ltd.
Gotenyama Hills
7-35, Kita-Shinagawa 4-chome
Shinagawa-Ku, Tokyo 140
Japan
Phone: 81-3-5488-3030
Fax: 81-3-5488-3031

**Mentor
Graphics®**



Printed on Recycled Paper

04-04-JC

TECH6300-w